



Publication Year	2022
Acceptance in OA	2025-03-05T13:18:12Z
Title	Highlights of the Square Kilometre Array Low Frequency (SKA-LOW) Telescope
Authors	LABATE, Maria Grazia, Waterson, Mark, Alachkar, Bassem, Hendre, Aniket, Lewis, Peter, BARTOLINI, MARCO, Dewdney, Peter
Publisher's version (DOI)	10.1117/1.JATIS.8.1.011024
Handle	http://hdl.handle.net/20.500.12386/36441
Journal	JOURNAL OF ASTRONOMICAL TELESCOPES, INSTRUMENTS, AND SYSTEMS
Volume	8

Highlights of the Square Kilometre Array Low Frequency (SKA-LOW) Telescope

Maria G. Labate,^{*} Mark Waterson, Bassem Alachkar[✉], Aniket Hendre, Peter Lewis, Marco Bartolini[✉], and Peter Dewdney

Square Kilometre Array Observatory, Jodrell Bank Observatory,
Lower Withington, United Kingdom

Abstract. The Square Kilometre Array (SKA) is an ambitious project to build the world's largest radio telescope to revolutionize our understanding of the Universe and the laws of fundamental physics. Geographically distributed between three host countries, and with more than a dozen member nations, the SKA is composed of two radio telescopes (SKA1-LOW and SKA1-MID) and a Headquarters facility. The SKA is now moving toward the start of the procurement phase and construction activities and is about to become a reality on the ground. Here, we focus on the SKA1-LOW and present the architectural highlights of what will be the most sensitive aperture array telescope on the earth, operating between 50 and 350 MHz. © 2022 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JATIS.8.1.011024](https://doi.org/10.1117/1.JATIS.8.1.011024)]

Keywords: radio astronomy; radio telescope; aperture arrays; phased arrays; Square Kilometre Array; interferometry.

Paper 21107SSV received Aug. 31, 2021; accepted for publication Mar. 14, 2022; published online Mar. 31, 2022.

1 Introduction

The Square Kilometre Array (SKA) Observatory will enable scientists to pursue world-leading scientific programs, and the SKA1 will be the first phase in the realization of this ambition.

The SKA Observatory functions as a single, integrated observatory distributed in three continents: Europe, Australia, and Africa. The Observatory is composed of a Global Headquarters in the United Kingdom and two SKA Telescopes, i.e., SKA1-LOW and SKA1-MID, located within radio quiet zones in Western Australia and South/Southern Africa, respectively (Fig. 1). The SKA1-LOW telescope will operate in the frequency range from 50 to 350 MHz, whereas the SKA1-MID telescope will provide continuous frequency coverage from 350 MHz through 15.35 GHz.¹

This paper aims to provide a general overview of the SKA1-LOW with a particular focus on the architectural highlights.

The paper is organized as follows. Section 2 describes the Science drivers for the SKA1-LOW. Section 3 explains SKA Observatory and SKA1-LOW locations. Section 4 shows Challenges that impact the SKA1-LOW architecture. Section 5 explains SKA1-LOW architecture. In Sec. 6, the Major components of the SKA1-LOW are described. Section 7 Concludes this paper.

2 Science Drivers for the SKA1-LOW

The aim of the SKA1 design is to conduct transformational science and to complement other front-line telescopes in the emerging era of multi-messenger astronomy. The SKA1 will contribute to the significant evolution of the various fields, such as Cosmology, Cradle of Life, Epoch of Reionization/Cosmic Dawn, Extragalactic Continuum, neutral hydrogen (HI) in Galaxies, Cosmic Magnetism, Our Galaxy, Pulsars, Solar Physics, and Transients. While the wealth of high-impact science that will benefit from the unprecedented sensitivity and characteristic of the SKA1-LOW is described in Ref. 2, a summary of goals is provided below:

^{*}Address all correspondence to Maria G. Labate, mariagrazia.labate@skao.int




Fig. 1 (a) The SKA headquarters in the United Kingdom, (b) two artistic impressions of the SKA1-LOW, and (c) SKA1-MID telescopes.

1. Penetration of the earliest stages of the Universe (Epoch of Reionization/Cosmic Dawn) as it transformed from a sea of neutral hydrogen to the first stars and galaxies.
2. Mapping the evolution of galaxies from their earliest formation until the current day.
3. Tracing the star-formation history of the Universe.
4. Discovering and timing new pulsars.
5. Advancing understanding of the source of fast and slow astronomical radio bursts.
6. Studying the mass composition of cosmic rays in the region of transition from Galactic to extragalactic origin.

The SKA1-LOW will prioritize imaging and spectral observations of the highly red-shifted 21-cm hyperfine line of HI from the Epoch of Reionization/Cosmic Dawn, as well as performing low radio frequency (RF) observations of pulsars (tending to have steep spectra), magnetized plasmas both in the Galaxy and intergalactic space, radio recombination lines, and potentially extrasolar planets. The SKA1-LOW discovery of additional pulsars will contribute to pulsar, gravitational wave, and interstellar medium studies. These investigations, along with unexpected discoveries that will follow, will revolutionize our understanding of the Universe.

3 SKA Observatory and SKA1-LOW Locations

Figure 2 shows the SKA Observatory, comprising the SKA1-LOW telescope, the SKA1-MID telescope, the Headquarters, and several SKA Regional Centers located off-site in partner countries. The geographic context of the SKA1-LOW is shown in Figs. 2 and 3.

1. The SKA Global Headquarters, already in place and operative in the United Kingdom, will eventually be home to more than 150 staff from more than 16 countries, tasked with the high-level planning and coordination for the SKA activities including management of the construction and remote monitoring of the operation of the SKA telescopes.
2. The observation proposal management and planning activities for the SKA1-LOW will commence at the SKA Global Headquarters and continue in the Science Operations Center (in Perth, Australia), where the SKA scientists and telescope operators will plan, execute, and monitor the observations, and perform quality control on the processed data via specialized tools.
3. The SKA1-LOW antenna array and primary electronics will be located on the Murchison Radio Observatory site at Boolardy, Western Australia, chosen after a competitive international site study and selection process, which evaluated several potential sites in terms of latitude, terrain, nearby population and potential radio frequency interference (RFI) sources, accessibility, and climate. After initial processing, digital data are transmitted by high-speed networks to a separate off-site facility [the Science Processing Center (SPC) in Perth] for further processing, quality control assessment and initial archiving.
4. Validated science data products will be distributed internationally to SKA Regional Centers, high-performance computing (HPC) facilities provided by SKA partner organizations, to further process, analyze, and store SKA data under the control of the user community science in the member countries.
5. The SKA1-LOW Engineering Operations Center (located at Geraldton, more than 300 km from the array location) will be the primary maintenance support facility, containing




Fig. 2 The SKA Observatory. The SKA facilities on site (i.e., Boolardy for SKA1-LOW) are shown in green, and the host-country off-site facilities are in purple. International entities are shown in pink.



Fig. 3 Geographic context for the planned locations of the SKA1-LOW array site, the EOC, the SPC and the SOC.

workshops, storage, and office space. The SKA1-LOW system Integration and Test Facility will also be in Geraldton, and it will support Integration, Verification, and Test activities to assure that telescope components perform properly.

4 Challenges that Impact the SKA1-LOW Architecture

4.1 Performance Objectives

The scientific performance is mainly determined by the following capabilities and performance characteristics, based on the science cases described in *Advancing Astrophysics with the Square Kilometre Array*² and the *SKA Phase 1 System Requirements Specification*³:

1. **Frequency Range:** the range of frequencies or wavelengths over which the telescope has significant sensitivity. For the SKA1-LOW, this is from 50 to 350 MHz, providing coverage of frequencies corresponding to red-shifted hydrogen in the early universe and the predicted Epoch of Reionization/Cosmic Dawn regime, and extending to meet the SKA1-MID range at 350 MHz.

2. Sensitivity: the ability to detect faint signals. The customary way to specify sensitivity, adopted by the SKA, is A_e/T_{sys} , where A_e is the effective collecting area including inefficiencies and losses, and T_{sys} is the total system noise, including sky noise and instrumental noise. The minimum SKA1-LOW sensitivity required at zenith, determined from high priority scientific requirements,² is from 68 to 625 m²/K depending on frequency (above ~500 m²/K at frequencies above 110 MHz).³ This is complementary to the sensitivity of the SKA-MID at higher frequencies and, when proposed, significantly exceeded the sensitivity of similar Low Frequency telescopes.
3. Bandwidth: the RF bandwidth that is available to the telescope at any one time. Sensitivity for wide-band (continuum) observations is proportional to the bandwidth. Although this does not confer additional sensitivity for spectral line observations, it does assist searches for spectral-line emission at unknown frequencies. The SKA1-LOW provides an instantaneous bandwidth of 300 MHz to enable time-domain de-dispersion of pulsar data over the widest available intervals.
4. Polarization capability: Orthogonal linear polarizations are sampled by the SKA1-LOW antennas providing the capability to measure and image polarization characteristics of radio emission.
5. Distribution of Collecting Area: at a given frequency, the sensitivity of the telescope to components of the spatial spectrum. The SKA1-LOW distribution of the collecting area (Sec. 5.1) is chosen in a compromise to support both sensitivity at maximum resolution, for precise imaging, and high surface-brightness sensitivity, for measurement of large-scale structures such as redshifted HI from the Epoch of Reionization/Cosmic Dawn.
6. Maximum Baseline: this determines the ultimate spatial resolution of the telescope. The resolution is determined approximately by the inverse of the maximum baseline (74 km for SKA1-LOW) and is ~10 arc sec at 110 MHz.
7. Temporal resolution: the ability to resolve temporal variations limited only by bandwidths and noise in both imaging (including integration periods as short as 0.25 s to support solar observations and rapidly moving objects) and pulsar timing (PST) to submillisecond resolution.
8. Specialized pulsar capabilities: Capabilities to both carry out searches for new pulsars over the entire visible sky, and extremely precise PST observations over a decade of elapsed time.
9. Very Long Baseline Interferometry: a capability to participate in observations with VLBI networks for which there is mutual sky visibility and frequency range compatibility. [Available telescopes are FAST, uGMRT, and VLBA antennas in Western United States and Hawaii, and antennas in Japan. In addition to what is currently available at this latitude and frequency range, the capability will allow to include suitable facilities that will become operational during the lifespan of the Observatory, e.g., the Commonwealth Scientific and Industrial Research Organisation (CSIRO) Bluering project in Australia.]
10. Processing capability along different dimensions: both spatial and spectral processing capabilities are present in the telescope as necessary to support imaging of the sky at a given frequency band in all four Stokes parameters ($IQUV$), as well as producing high-resolution spectra over a defined area of sky.
11. Observing flexibility and response time: Flexibility in the use of capabilities and resource management to help to optimize observations while dealing with uncertainties of observing demands and conditions, and with the discovery of the unknown. To take advantage of the aperture array technology while keeping within the cost and complexity constraints, SKA1-LOW operational capabilities allow, e.g., to:
 - a. Divide the telescope array into subarrays (i.e., allocate and configure a subset of telescope resources—antennas, beamformers, correlators, and processing resources—just sufficient for the observation).
 - b. Support multiple scientific uses of the same observations from the same subarrays through commensal data collection, e.g., combining pulsar searches and imaging observations.

- c. Form multiple beams within multiple subarrays.
- d. Detect and respond to fast (ms) astronomical transient events (e.g., fast radio bursts).
- e. Form images on multi-second time scales (slow astronomical events).
- f. Rapidly change observing programs as a result of an external or internal trigger (e.g., a supernova event detected by another telescope).

The achievement of the demanding scientific goals, while dealing with cost, time, physical, and environmental constraints, represents a huge unique challenge. Because of its unique combination of resolution, survey speed and sensitivity, SKA1-LOW will represent a significant improvement over existing radio telescopes in this frequency range, as discussed in Sec. 4.2.

4.2 Comparison with State-of-the-Art Telescopes on the SKA1-LOW Frequency Range

The SKA project represents the first opportunity in a decade to build a large new radio telescope in the dm/cm wavelength range, and the first large expansion in the meter wavelength range. As illustrated by the two-volume exploration of SKA science,² more than just large collecting area and low-noise designs are required.

As the next generation telescope, the SKA's astronomical performance must be a significant step over currently available telescopes. A driving concept for the SKA has been sensitivity, and a signature requirement and basic assumption behind the most important SKA observations is to be able to integrate for at least 1000 h, limited in sensitivity only by uncontrollable, natural noise over the full field-of-view at the highest spatial resolution.

While many current telescopes overlap with some of the capabilities of the SKA1-LOW, the most comparable designs are the LOW-Frequency ARray (LOFAR),⁴ Long Wavelength Array (LWA),^{5,6} and Murchison Widefield Array (MWA)⁷ radio-telescopes.

The SKA1-LOW design covers a similar frequency to these telescopes but provides an overall sensitivity increase of more than an order of magnitude (Fig. 4) while being optimized for brightness temperature sensitivity by arranging much of the collecting area in a very compact array. In addition to traditional imaging, the SKA1-LOW will support commensal time-domain processing for pulsar and transient observations. In summary, the SKA1-LOW design outlined in this document will be a major step forward in astronomical performance, and its location in the Southern hemisphere will complement similar telescopes in the North, as well as the very large optical/infrared telescopes and Atacama Large Millimeter Array (ALMA) in the South. MWA




Fig. 4 The predicted sensitivity of SKA1, compared to indicative sensitivity performance of some other existing and planned radio telescopes covering a similar frequency range.⁸ (Note: this figure is taken from Ref. 1 and does not include planned facilities, such as, e.g., LOFAR2.0.)

and LOFAR are both embarking on significant upgrades,⁹ however, improvements on flexibility and correlator capabilities are emphasized rather than significant increases in collecting area. Further information on performance parameters of the SKA and key existing facilities can be found in Refs. 1, 10, and 11.

4.3 Most Significant Technical Challenges

Some of the technical challenges that drive the SKA1-LOW design are:

1. Very high number of antennas: SKA1-LOW's sensitivity goals imply a large collecting area, which at these frequencies implies a telescope with the largest in terms of the number of antennas and a concomitant facility to process the approximately quarter million wide-band signals emanating from the antennas. The SKA1-LOW provides a feasible, robust, and reliable processing system with many simple signal chains combined with a maximally parallel signal processing design.
2. Station aperture sampling: the aperture array concept used for a very wide frequency range implies a fundamental compromise between frequency range and aperture sampling. For maximum aperture efficiency, the antennas would be spaced one half wavelength apart to achieve the maximum number of independent samples of the incoming wave. However, this would result in a very regular structure to the layout (leading to grating lobes and scan-blindness), high levels of mutual coupling between antennas, and need for many more antennas than can be reasonably afforded to cover the desired aperture area. Conversely, spacing the antennas at the optimum distance for the lowest frequency would result in an extremely sparse sampling at the highest, significantly reducing aperture efficiency at these frequencies. A simple log-periodic antenna designed to cover the frequency range would have elements with a maximum dimension of the order of $\lambda/2 \times \lambda/2$ at the lowest frequency, where λ is the observing wavelength. This antenna would be too big to fit within the physical aperture; as a result the SKA1-LOW design utilizes a log-periodic antenna structure with shaped dipole elements to realize a compact design, which is physically smaller than its apparent electrical length at the lowest frequency to counteract this effect. As the frequency increases, the individual antenna locations are farther apart in wavelengths, eventually reaching inter-element distances of more than $\lambda/2$. At this distance, the array is no longer Nyquist sampled. The frequency at which this happens is called the sparse-dense transition, which for the SKA1-LOW design is ~ 94 MHz. At these frequencies, the collecting area is approximately the physical area. At higher frequencies, the collecting area decreases as λ^2 . Above the sparse-dense transition frequency, the station beam will contain far-out sidelobes randomly placed on the sky (Sec. 5.1), limited in extent only by the beamwidth of the individual antennas.
3. Mutual Coupling between antenna elements: the adjacent antenna elements are strongly coupled at frequencies below the sparse-dense transition. This results in performance and calibration challenges that are mitigated through randomizing the antenna layouts, comprehensive electromagnetic modeling of the resulting embedded antenna element patterns, and advanced calibration strategies.
4. Long baselines: the largest dimension of the array layout results in the outer stations being in excess of 40 km from the central area of the array, with consequent challenges in terms of maintenance, logistics, and high cost of the power and signal distribution (PaSD) over such a large area. The SKA1-LOW addresses these challenges through an architecture providing the outer stations of a proper infrastructure, intermediate processing, remote photovoltaic power generation,¹ and optimal long-distance distribution of synchronization and timing (SAT) signals. Differences in the ground characteristics/properties, topology, and aspects related to earth curvature, as well as the impact of ionospheric effects over such a wide area, have also been carefully considered through, e.g., the use of the antenna ground plane and proper calibration strategies.¹²
5. Remote location: SKA1-LOW will be built in a very remote part of the world, the Murchison region of Western Australia. Remoteness helps to minimize RFI that

interferes with the astronomical signals but adds another layer of complexity to its delivery. The SKA1-LOW distributed system architecture is such that fewer operational staff need to live in a remote area. Also, each component must be qualified to withstand the tough environmental conditions on the site.

6. High real-time data rates: the SKA1-LOW front-end will produce up to 10 Tb/s of data, which must be reduced through real-time data processing and analysis before being made available to the global scientific community. The distributed computing and data networking are major drivers of the overall observatory architecture.
7. Real-time high-performance science computing: the volume of data ingested for science processing is so high that it cannot all realistically be stored. This means that unlike most radio telescopes, SKA1-LOW will require real-time science data processing (SDP), with only a subset of raw data being archived and distributed for scientific use.
8. Signal sampling coherence and time-keeping stability: Maintaining coherence between all the RF inputs is at the core of every aperture synthesis telescope. Simply put, RF inputs must be synchronized to $\sim 1^\circ$ of phase at the highest frequency corresponding to 8 ps at 350 MHz. Moreover, to meet requirements for time-domain astronomy, SKA1-LOW must maintain a stable time reference over a period of at least a decade. Coordinated Universal Time (UTC) is the globally recognized standard for time, and SKA1-LOW must be able to trace the occurrence of events (e.g., pulsar times-of-arrival and transient signals) to UTC with an accuracy of 9 ns (1-sigma) over 10 years. The SKA1-LOW provides up-to-date methods of keeping accurate time on site and distributing the timing and clock signals over long distances in the face of fluctuating environmental conditions.^{13,14}
9. High operational availability: Achieving the challenging scientific goals of the project in a cost-constrained operational environment requires a system design that is reliable and robust to failure, without relying on costly levels of redundancy, to maintain an Operational Availability level of 95% (excluding weather issues).³ SKA1-LOW has $\sim 130,000$ antennas, 260,000 signal chains, long-distance data transport, and many other high-volume or expensive features. Maintaining high system availability requires explicit engineering design in the Reliability, Availability, and Maintainability areas. Even the simplest devices deployed at high volume in the site's harsh desert environment can fail frequently. The SKA1-LOW architecture combines high reliability with a measure of redundancy and allow ongoing maintenance in tandem with operations.
10. Electromagnetic Interference (EMI) emissions: in addition to the challenge of RFI from external sources, behavior, and its control of the devices (including single antennas) outside the processing facilities, as well as the computing part, data transmission and human activity must also be designed to avoid self-interference (EMI), driving the need to keep computing and people away from the array. An SKA standard has been put in place to ensure that internally generated RFI is controlled.¹⁵
11. Long-term observatory vision: Historical evidence provided by generations of radio telescopes indicates that its most important discoveries in its anticipated 50-year lifespan cannot be envisaged. The SKA1-LOW architecture is designed to be, as much as possible, modular, flexible, and expandable with minimal cost and disruption. This is particularly true for computing and network hardware, which must be able to grow to meet new scientific opportunities and expanding demand for more science products.

5 SKA1-LOW Architecture

The station design and parameters have evolved during nearly 20 years of studies and precursor telescope projects to obtain the best tradeoff in terms of performance and cost for the different science cases and for calibration purposes. The key advance of the SKA1-LOW design, beyond the previous generation of large-scale, low-frequency radio telescopes, is the adoption of direct digitization of all the primary antenna elements at baseband. This enables completely digital channelization and beamforming, programmable aperture weighting and substation apertures, frequency channel duplication or re-use, and commensal time-domain processing for pulsar and




Fig. 5 (a) Antenna, (b) field station of 256 antennas, and (c) SKA1-LOW array of 512 field stations. While the SKA1-LOW is an artistic impression, both the antenna and the field stations are two prototypes part of the aperture array verification system (AAVS) 2.0 demonstrator at the Murchison site. (A technology demonstrator platform called the AAVS was built at the Murchison site in Western Australia to gain field experience under true site conditions and increase confidence on the station performance and calibration aspects. A series of iterations of the system have been built, over the past 4 years, culminating in the AAVS2.0 system now in use for developing calibration strategies.^{8,11,15–24}) (d) Conceptual illustration of the different beams, from left to right: antenna beam, substation beam generated by a subgroup of antennas in the station, station beam, and tied-array beam summed from several stations.

transient detection. The resulting digitally steered aperture array is composed of 131,072 dual-polarized antennas grouped in 512 stations of 256 antennas (Fig. 5). The RF system uses an FX correlator architecture with a two-stage frequency channelization scheme. Each station beam-former filters the full 50 to 350 MHz band into 384 oversampled coarse channels (resulting in no inter-channel loss) and transmits up to 48 independent sets of station beams (totaling 300 MHz) to the Correlator and (Array) Beamformer. Further channelization into finer channels of programmable resolution, down to 250 Hz in optional zoom modes, precedes correlation into visibility products. Simultaneous time-domain processing of tied-array beams supports pulsar search and timing pipelines.³

SKA1-LOW operates concurrently in imaging mode and non-imaging mode with operation of between 1 and 16 subarrays. Each subarray is programmable as a separate conceptual telescope in terms of antenna pointing, band selection and the setting of configurable imaging and non-imaging parameters. Within any subarray, the imaging correlator and pulsar (array) beam-formers must operate concurrently, and therefore, separate correlator, Pulsar Search, and Pulsar Timing channelizers are required simultaneously.

5.1 SKA1-LOW Configuration

Part of the telescope design process is to ensure that astronomical performance will be a major improvement over currently available telescopes. The array configuration (i.e., location of the different antennas on site) has been chosen to maximize overall scientific performance while taking into account local conditions (topography, access to land, culturally sensitive areas, proximity to power, and local sources of RFI).

Guided by astronomy, the design principles of the array configuration are:

1. Sensitivity: very high brightness temperature sensitivity over the range of spatial scales is required for the identified science cases. Sensitivity has been the driving concept for the SKA alongside appropriate resolution and dynamic range to achieve a major step in collecting.
2. Resolution: sufficient resolution is required to accurately map the point-source distribution at low frequencies.^{2,11,16} This will be useful for general astronomy applications, but crucial in removing the so-called “foreground” in Epoch of Reionization/Cosmic Dawn observations. Resolution is almost as important as sensitivity and for many types of observations (e.g., studies of weak and distant radio sources) resolution and sensitivity must be matched.¹ Except for very long integration times, SKA1-LOW continuum surveys are not seriously impacted by confusion noise down to a frequency of ~ 110 MHz. (Note that narrow spectral line observations will never be confusion-limited.)¹
3. Snapshot coverage: the two-dimensional (2D) distribution of baseline vectors determines the instantaneous (“snapshot”) spatial frequency (uv plane) coverage (Fig. 7). This will also be important for removing foregrounds and for solving the ionospheric phase screen over the array in near-real-time.
4. Long-track coverage: the 2D distribution also determines the long-track coverage of the uv plane in rotation-synthesis observations. The amplitude response of the individual log-periodic antenna elements falls off gradually with zenith angle, and projection effects reduce the quality and sensitivity of the beam at low elevations. This limits high quality observations of a single field to within 4 h of the meridian (although the system will acquire data from any pointing angle).
5. Scale-free: scale-free distributions of baselines ensure that the telescope does not favor a particular spatial scale, apart from that defined by the core itself.
6. Synthesized Point Spread Function: high quality synthesized beam (Point Spread Function with low sidelobes)¹¹ is required. The response to sources far away from the field center is the product of the station beam and the beam of the individual antenna elements.
7. Beam and the synthesized beam: because the station apertures are under-sampled at high frequencies, there will be sidelobes of the station beam that may include strong sources.
8. Ionospheric Phase Screen: the ionosphere is a critical factor at these radio frequencies, and the physical extent of the array is larger than the scale of ionospheric distortion features. The configuration of the outer stations is important to provide information needed to characterize the structure of the ionospheric phase screen, over both the core and the rest of the array.

These principles have led to the SKA1-LOW configuration shown in Figs. 6 and 8.




Fig. 6 The geographic location of SKA1-LOW in Western Australia. The black dots are the locations of the field stations (in the core) and clusters (in the spiral arms), and the blue outline shows the approximate boundary of the currently acquired property.

To simultaneously reach resolution and sensitivity goals, the SKA1 telescopes use the aperture synthesis concept in which multiple connected antennas are combined to form a single large collecting aperture. As shown in Fig. 6, the SKA1-LOW is made of 512 circular aperture arrays (“Field Stations,” Fig. 5) of 256 antennas each. Each field station has a diameter of about 40 m and its size has been chosen based on the need to achieve a field of view large enough to sample significant volumes of the Universe for the Epoch of Reionization/Cosmic Dawn 21 cm experiment^{2,11,16} while maintaining sufficient sensitivity for station calibration.

The antenna positions within each circular Field Station are randomized to avoid grating lobes, sidelobe structure, and scan-blindness effects. Internal studies have shown that for this choice of element density and station size, of order eight unique randomized layouts are sufficient to mitigate residual sidelobe structure in the correlated array beam without the need for rotating stations with respect to one another.

Roughly half of the effective collecting area is condensed in a central core area where 224 Field Stations are arranged with radially decreasing density to provide enhanced sensitivity to extended regions of low brightness and a sensitive subarray for pulsar and transient astronomy. The remaining 288 Field Stations are placed along three pseudo-spiral arms (extending ~ 35 km from the array center) to provide a logarithmic decrease in collecting area with radius beyond the 1-km core to ensure the uv plane¹ coverage and avoid favoring any particular spatial scale.

The Field Stations in the core are randomly (three radial zones were defined, each with a slowly declining density of stations that needed to be placed in each. Locations for station placement were chosen at random (using a random number generator for a radius and an azimuth) to be placed within each radial zone starting from the center with the only constraint that they not lie on top of any previously placed station. Subsequently, some stations were shifted due to cultural heritage constraints arranged across the entire core area, whereas the spiral arms are populated with 16 clusters of six Field Stations arranged randomly over an area 100 to 150 m in diameter. This arrangement has been chosen to provide high sensitivity (more collecting area) in the outer part of the array representing a cost-effective approach that reduces uv coverage only slightly. To improve uv coverage in snapshot imaging and fit within the boundaries of the observatory property, three sets of clusters at the same radius in each arm have been rotated by 10 deg to 20 deg²⁰ with respect to a regular spiral. This does not change the radial distribution of collecting area but fills in missing baseline vectors in the uv plane.

The uv coverage resulting from this configuration is well sampled for snapshot observations, and as can be seen in Fig. 7 fills in well over a longer tracked observation. It should be noted that track lengths are limited in an aperture array of this type due to beam pattern evolution and projection effects at lower elevation angles. The SKA1-LOW precision performance parameters are therefore only specified to zenith angles of ± 30 deg although the system will operate to any desired pointing angle (with reduced performance).




Fig. 7 (a) Snapshot and (b) 4-h track visibility sampling at a single frequency for a source at declination -30 deg.




Fig. 8 (a) Full array configuration. Each dot is a cluster of six stations. (b) Inner 3.4-km diameter array configuration. Each dot is a station. (c) Station beamformers, and the correlator and (array) beamformer and pulsar processors, which transmit output data to the SDPs and VLBI terminals in Perth.

5.2 SKA1-LOW Processing Facilities arrangement

To mitigate some of the challenges of building and operating the telescope in such a remote, hostile environment, the SKA1-LOW signal processing and support equipment will be installed within thermally controlled and EMI-shielded processing facilities to centralize maintenance, support, and to control radiated emissions. Signals from the antennas in the Field Stations are transmitted to the nearest processing facilities through RF-over-Fiber (RFoF) analog links²¹ (Fig. 8), which enable concentrating the signal processing equipment for multiple stations in a small number of locations: 296 stations in the inner area of the array will be serviced by

a large, well shielded data-center building [the Central Processing Facility (CPF)], whereas the remaining 216 stations along the outer spiral arms are clustered into groups of 6, each served by a smaller shielded building located nearby [Remote Processing Facility (RPF)]. This simplifies providing infrastructure services (power, cooling, and network connectivity), reduces the number of maintenance sites, and enables designing the digital signal processing (DSP) electronics to operate in a normal server-center environment.

In particular, the large CPF will house the digitization, timing, and communications equipment supporting the (296) stations within a central area of 3.4-km diameter (i.e., the 1-km core and the most central four clusters in each arm). The CPF is located ~ 400 m from the core center. This ensures that the RFoF links are no longer than 4 km but are far enough away that residual radio emissions do not affect the array.

As the distance from the outer locations to the central area exceeds the distance over which simple RFoF links can transmit analog data without degradation,²¹ the smaller RPFs are provided to house the digitizers for these antennas. In particular, each of the 36 (12" each arm) outer clusters is served by one RPF located within 575 m of the stations making up each cluster. The RPFs are linked to the CPF via digital optical fiber links to communicate science, timing synchronization, monitoring, and control data.

Central and RPFs house equipment needed to support processing of logical station data, including signal conditioning, digitization, station beamforming, networking, power, and local control. The CPF will also host the station control and calibration computing, the Correlator and (Array) Beamformer, the pulsar search and PST processors, the SKA1-LOW timescale reference system, and all the necessary infrastructure to support these functions [After the submission of this paper, an Engineering Change Proposal (ECP), proposing to move the Correlator and (Array) Beamformer, the Pulsar Search, and the PST subsystems from the CPF to the SPC, was accepted. The architecture, hardware, firmware, and software of these subsystems are not affected, and the change will mainly affect the networking and infrastructure. The implementation of this ECP will be covered in future SKA1-LOW architectural papers.]. All these facilities will be linked by roads and data transmission fibers. Power will be supplied by a distributed system that includes significant renewable power elements.

5.3 SKA1-LOW Functional Description

Figure 9 shows the major components and the signal flow through the system:

1. In each Field Node, the signals received at each antenna element are amplified by a low-noise amplifier (LNA), conditioned, converted into analog optical signals and transmitted, via optical RFoF links (i.e., Fiber Cable Assembly), to nearby processing facilities from hundreds of meters to a few kilometers away.
2. Digitization is carried out in the signal processing subsystems located inside the Central and RPFs, where the first stage of beamforming (station beamforming) takes place. The channelized, beamformed data streams from the station beamformers are sent to the Correlator and (Array) Beamformer for fine-channel filtering and correlation.
3. A transient capture buffer is also implemented by forwarding re-quantized station beams to a large ring buffer memory in the Monitor Control and Calibration Subsystem computing cluster located in the CPF. Trigger signals from several sources, including real-time transient detection software and external event notifications, freeze the contents of the buffer for download to the SPC for in-depth transient analysis.
4. The Correlator and (Array) Beamformer channelizes the signals into finer channels and then correlates the signals from each channelized station beam (or substation beam) to produce output visibilities, which are passed on to the SDP. The Correlator and (array) Beamformer also coherently adds signals from all stations together to form tied-array beams to be used for Pulsar Search, Pulsar Timing, and VLBI. Delays are implemented as sample delays and linear phase-slopes across the frequency band. Each coherent beam on a 20-km aperture (For pulsar studies an array beamformer and a centrally condensed core area configuration are used, as angular resolution is not required, and as computing costs/requirements scale strongly with the filling factor of the array.) has a Half-Power




Fig. 9 SKA1-LOW architecture.

Beam Width (HPBW) of $\sim \lambda / (D\text{-array})$, where λ is the wavelength, and $D\text{-array}$ is the diameter of the array, which is limited to 20 km for pulsar processing. The HPBW of the station beam is $\sim \lambda / (D\text{-station})$, where $D\text{-station}$ is the diameter of the station, which is about 40 m. To fill the primary beam of the station, 500 Pulsar Search beams are formed. PST beams are far fewer than the Pulsar Search beams, as they are used for targeted observations of known pulsars exploiting the full system bandwidth of 300 MHz for higher sensitivity.

5. The pulsar search subsystem (PSS) provides commensal processing of up to 500 tied-array beams to extract potential pulsar signals from these tied-array beams. The PSS outputs pulsar candidates as a consolidated dataset containing the prospective pulsar parameters, which is passed to the SDP for further verification and analysis. The PSS can also generate a list of single-pulse candidates for further analyses by the SDP. Real-time processing of these candidates will be used to trigger the monitor control and calibration subsystem to freeze and save sections of the transient capture buffer. Similarly, the PST subsystem is a smaller DSP subsystem designed to measure the times-of-arrival of pulses from known pulsars very accurately, along with other parameters (pulse rate, rate of change, pulse profile, etc.).
6. The CPF also houses the SKA1-LOW SAT system, responsible for providing SAT signals to sample the astronomical data coherently captured by each receptor and tie it to the time as accurately as possible. The SAT system is composed of the Observatory Clock (time-scale) and the Frequency and Timing Distribution subsystems. The first generates high accuracy and stable time and frequency signals, whereas the latter is responsible for distributing these to the remote locations. The timescale is designed to participate in the group of organizations coordinated by Bureau International des Poids et Mesures (BIPM) for calculation of the UTC, the internationally recognized time standard. The BIPM is the organization responsible for coordinating and disseminating global estimates of atomic time as UTC. The SKA1-LOW timescale is synchronized with UTC via satellite through a set of Global Navigation Satellite System (GNSS) receivers. Additional hardware is used to distribute reference signals generated by timescale to subsystems that need it.
7. The Observation Monitoring and Control (OMC) system includes two main subsystems: the Telescope Monitoring and Control (TMC) and the Observatory Science Operations (OSO). The TMC is the control and monitoring system for SKA1-LOW. It orchestrates the hardware and software systems and facilitates maintenance of telescope systems




Fig. 10 Approximate data rates between different SKA1-LOW components and the SKA regional centers.

through logging of “health” parameters and provides support for performing diagnostics based on archived data. It delivers these data to operators, maintainers, engineers, and science users through user interfaces. The OSOs subsystem manages the observation life-cycle, from proposal submission to observation execution.

8. The SDP at the SPC receives visibilities from the Correlator and (array) Beamformer through an ingest switch, which directs data to a highly parallelized supercomputer, which carries out calibration and imaging. The resulting multi-dimensional images (two spatial, one frequency, and four polarization parameters) are stored in the Science Data Archive, and ultimately distributed by long-haul fibers to the SKA regional centers.
9. VLBI data can be streamed in real-time to a VLBI terminal, which can transmit data in standard form to a remote VLBI correlator (e-VLBI mode) or can record data for later streaming after the observation has been performed.

Figure 9 shows an outline of data flow through the major components of the system, whereas Fig. 10 shows the data rates between them.

6 Major SKA1-LOW Components

The SKA1-LOW architecture overview showing the major subsystems and the signal flow through the system was shown (Fig. 9) and introduced in Sec. 5.3. This section aims to provide further details on the major components, i.e.,

1. *Field Node*: signal capturing and front-end components.
2. *Signal Processing Subsystem*: station beamforming.
3. *Correlator & Array Beamformer*: generation of visibilities and tied-array beamforming.
4. *Pulsar Search* subsystem: processing for pulsar and fast transient search.
5. *Pulsar Timing* subsystem: precisely measure of the times of arrival (TOA) of known pulsars.
6. *Science Data Processor*: processing of observed data into science data product.
7. *Observation Monitoring and Control*: managing observation planning, scheduling, and execution.
8. *Synchronization & Timing system*: providing SAT signals.
9. *Networks*: providing signal and data transport infrastructure.

6.1 Field Station and Field Node

The performance of the entire telescope is largely constrained by the design of the 512 field nodes that comprise the entire SKA1-LOW telescope array. These consist of

1. A 256 antenna-element array, configured as a quasi-circular array (field station), to capture and amplify the signal. The log-periodic dipole antennas are deployed over a 42-m diameter ground plane, consisting of a metal mesh over bare ground.
2. A PaSD system, made of 24 SMART Boxes and a Field Node Distribution Hub (FNDH), to power, monitor, control, and collect the signal from the antennas.
3. Fiber cable assemblies (to carry the signals from the stations to the processing facilities).

There are very few antenna-element designs that can meet the frequency range requirement. Log-periodic dipole antennas belong to a family of self-similar designs whose most important feature is wide bandwidth ratio (ratio of upper to lower frequency limits). They consist of two orthogonal sets of half-wavelength dipoles of different dimensions spaced along a support structure and fed by two transmission lines, also known as booms (Fig. 12). The useable frequency range is determined by the ratio of the lengths of the longest to shortest dipoles. This feature enables SKA1-LOW frequency coverage from 50 to 350 MHz, albeit with a very sparse array at the high end. Two LNAs, one for each orthogonal polarization, are attached to the transmission line at the top of the structure.

The design and optimization of the SKA1-LOW antenna has taken several years and in the course of the project several versions have been proposed.^{1,23–25} The version selected as reference antenna is the so-called SKALA4.1. The main characteristics of the SKALA4.1 log-periodic antenna are:³

1. 20 dipole elements: 19 triangular-tooth plus one electrically short bow-tie at the bottom of the antenna. A standard log-periodic antenna is composed by a series of half-wave dipoles. In this case however, the longest dipoles have been shortened to permit a closer packing of the antenna elements to raise the sparse-density transition frequency (i.e., improve the spatial sampling of the aperture at high frequencies). In fact, while the minimum operative frequency of SKA1-LOW would require a dipole length of 3 m (corresponding to $\lambda/2$ at 50 MHz), the required maximum dimension of the antenna limited this length to 1.6 m. However, the bow-tie shape of the bottom dipole²³ aims to compensate for the reduction in length by allowing it to slightly move down in frequency its resonance and increasing the bandwidth. Moreover, the slight reduction on matching, directivity, and smoothness with respect to a nominal $\lambda/2$ dipole are compensated by the fact that, at those low frequencies, where already the array is densely packed and mutual coupling effects stronger, the system noise is determined by sky noise and not instrumental noise.²
2. Solid dipoles on the high-frequency elements and wire dipoles on the low-frequency ones.
3. The angle between the two halves of the transmission line is one degree, which improves the matching and smoothness of the frequency response.
4. Electrical connection of the antenna to the ground plane, which discharges static buildup, common in dusty environments.
5. Each polarization is matched to a single-ended 50- Ω LNA.

The purpose of the mesh ground plane is to provide an electrically stable reference plane for the antenna elements. Simulations have shown that without a ground plane, changing ground conditions affect the antenna response, especially after rain. The ground plane consists of 214 steel mesh panels (2.4×3 m) that are electrically bonded. The 4-mm mesh wires are spaced in a 50×50 mm grid. The grid also provides a means to easily attach the antennas.

The outputs from the LNAs are connected via two coaxial cables to Front-End Modules, which are housed in one of 24 locally situated SMART Boxes (Figs. 11 and 12). The positions of the SMART Boxes are arranged within the ground plane diameter to ensure that the coaxial cable has a maximum length <10 m. This length is short enough that residual amplitude structure in the passband due to standing waves is minimized within the delay range window of a potential Epoch of Reionization/Cosmic Dawn signal. Matching circuits are provided at each end of these cables to minimize reflections.

Each Front-End Module contains an RF amplifier, high-pass out-of-band RF filter, and RFoF optical laser drive module, which transmits the two RF signals at different optical wavelengths.²¹




Fig. 11 (top) Drawing of a field station with his PaSD.




Fig. 12 (a) A drawing of the SKALA4.1 antenna element. The LNA is placed at the top of the antenna. The central support is also the transmission line. Note the two perpendicular planes of polarization. (b) Prototypes of the SKALA4.1 antenna and SMART Boxes used as part of the AAVS2.²²

Both the coaxial and the RFoF cables within the Field Station are exposed on top of the ground plane for deployment and maintenance reasons.

At the FNDH (enclosure and primary control node of the Field Station), the fibers carrying analog signals from each antenna and each polarization are spliced into larger trunk cables (288-core Fiber Cable Assembly). To provide protection and thermal stability, the Fiber Cable Assembly cable runs in trenches alongside the station power cables back to the Central or RPFs. The use of RFoF technology allows, together with other benefits,²¹ to cover the long distances required by the SKA1-LOW topology, which would otherwise lead to unacceptable levels of attenuation.²¹ Each Fiber Cable Assembly cable passes into the shielded area of the processing facility through a waveguide feedthrough and terminates in a rack containing 16 digitizer modules [Tile Processing Modules (TPMs)], which convert the optical signals back to electrical RF for further processing.

Within the FNDH, a power control and conversion module provides monitor and control capabilities via a bidirectional link to the Monitor Control and Calibration Subsystem, and mains AC power is down-converted and distributed to the SMART Boxes along with LMC signals. The LNA bias power is then distributed to each antenna via coaxial cables.

6.2 Station Signal Processing Subsystem

Figures 13 and 14 show the station processing architecture. Each station, consisting of 16 TPMs, implements this functionality in parallel, while the beams are formed by a process-and-forward partial-beam architecture.

1. Fibers, carrying the analog signals from each antenna in the station, enter the inputs of 16 TPMs²⁶ where conversion to electrical signals, amplification, band-pass filtering and digitization occur. Each TPM handles 16 dual-polarized RF signals. The TPM racks are housed in the RF-shielded, controlled environments of the Central or RPFs.
2. In each TPM the antenna signals are digitized with a 800-MHz sampling clock at 12 bits. The RFI environment at the site creates an effective signal dynamic range that significantly exceeds the total sky power integrated across the bandwidth of the receiver. To avoid signal chain complexity it was determined that, for cost, power, and performance reasons, a high-performance 12 bit digitizer provided the best solution combining an efficient data interface with sufficient dynamic range. Only the 8 most significant bits from each digitizer are transferred to a Polyphase Filter Bank (PFB), which produces 512 coarse




Fig. 13 A block diagram of one station, including the functional scheme of a TPM: analog section on the left (PreADU), digital PB on the right (ADU), digital flow from left (antennas) to right [outputs to the correlator and (array) beamformer, transient buffer, and station monitor control and calibration subsystem].




Fig. 14 Outline processing flow for tile and station beamforming.

channels covering the entire frequency range up to the Nyquist sampling frequency; 384 frequency channels cover the 300-MHz science bandwidth in overlapping 781 kHz channels while the remaining channels are available for diagnostics and monitoring of out-of-band signals.

3. Calibration, pointing, and beamforming operations are carried out independently on each of the resulting 781.25 kHz channel data streams. A static delay term (in samples) is applied to compensate for the RFoF and signal chain electrical length and an adjustable gain equalization is applied. At each observation scan boundary, calibration of the chain can be adjusted, if necessary, by updating the complex calibration coefficients for each stream prior to beamforming. Studies have shown that system stability is sufficient to avoid the need for real-time computation of calibration coefficients, therefore calibration coefficients will be obtained from a database of pre-validated parameters selected according to current environmental conditions. Calibration coefficients will be derived from on-sky measurements processed offline to generate updated coefficients correcting each channel's complex receiver gain, antenna gain, polarization correction, and a bandpass equalization term to remove slopes in signal amplitude across the received band. Calibration processing is implemented in the Monitor Control and Calibration Subsystem^{1,27} software²⁸ running on a dedicated computer cluster at the CPF. Pointing, forming a beam in the required direction on the sky, is accomplished by calculating an individual delay for each antenna location from the commanded pointing coordinates. These delays are converted to per-channel phase terms in the TPM beamformer firmware and added to the calibration phase term mentioned above. Tracking observations are performed by locally incrementing the channel delays by a delay-rate term; the pointing coefficients (delay and delay-rate) are updated as often as required to maintain the beam center on the desired sky point. Non-sidereal tracking rates (e.g., Sun, planets, and LEO objects) are supported.
4. After calibration and pointing terms are applied, the data are rescaled to preserve dynamic range, and the 16-antenna channelized sample streams for each channel are summed to form a partial station beam and forwarded to the next node in the chain. The partial beams from the 16 other TPMs making up the station are summed in series to form a full station beam. To support maximal scientific flexibility and observing efficiency while limiting control complexity, the system processing capacity of 300 MHz (384 channels) can be managed as up to 48 independent station beams in groups of 8 channels, where individual frequency channels can be duplicated or omitted as needed. In a special holography mode, 384 single-channel beams can be formed to support beam-shape mapping.
5. The TPMs are all attached to a 40 Gb network "bus" within each station from which the TPMs can read or write packetized data. In this way a chain is formed to sum together all the partial beams. The final TPM in each chain forwards its station-beam stream to the Correlator and (Array) Beamformer subsystem over a 100 Gb link.

A more detailed description of the station processing, including details on the steps shown in Fig. 14, can be found in Ref. 26.

6.3 Correlator and Array Beamformer

The Correlator and (Array) Beamformer²⁹ receives, from the Signal Processing Subsystem, station-beam data oversampled by a ratio of $32/27 = 1.19$, resulting in overlapping channels to avoid frequency aliasing near the channel edges when this data is processed. The total rate of the data received from the 512 stations, 300 MHz bandwidth, dual-polarization signals, and two 8-bits words for each complex sample, is 5.8 Tbps (Fig. 10).

The correlator in the Correlator and (Array) Beamformer subsystem calculates cross-correlation products between the stations in the array, and autocorrelations as well. The DSP produces visibilities at a 5.4-kHz resolution (55 k channels), in the normal imaging mode. Each complex product has a minimum specified integration time of 0.9 s; therefore, the maximum visibility data rate is 29 Giga-visibilities/s. Each visibility is defined as 64-bit complex data, with additional 8 bits for the time centroid index and 8 bits for a quality value (i.e., the fraction of data that was correlated). At 80 bits per visibility, the rate of data sent to the SDP is ~ 2.3 Tbps. In parallel

with the normal imaging mode, a processing of 64 zoom bands can be conducted (imaging zoom mode). The zoom windows have a bandwidth ranging from 4 to 256 MHz and can be placed anywhere in the 300 MHz.

Instead of the zoom mode processing, the Correlator and (Array) Beamformer can be used in pulsar-search and pulsar-timing modes. The tied-array beamformer, included in the Correlator and (Array) Beamformer subsystem, adds coherently the allocated stations to form power beams for Pulsar Search, and voltage beams for PST and VLBI. Up to 500 Pulsar Search beams can be formed using nearly 15-kHz wide frequency channels per beam. The output data rate, for 500 beams, 2 polarizations, 118-MHz bandwidth, and two 8-bits words per sample, is 1.9 Tbps. In PST, more 16 full bandwidth beams are available for targeted observations of known pulsars. The rate of the output data of 16 beams, two polarizations, 300-MHz bandwidth, and two 8-bit words for each complex sample, is about 0.15 Tbps. Four PST beams can be reassigned to produce VLBI beams. Each VLBI beam has up to four VLBI standard sample rate subbands, each with a bandwidth of between 1 and 64 MHz.

For correlation and beamforming, the signals are channelized in the frequency domain by channelizers. The correlator is of type FX, which is based on first performing frequency decomposition to frequency channels, then correlating the signals per frequency channel.

The architecture of the Correlator and (Array) Beamformer system can be described as composed of three major subassemblies: a processing subassembly, a data routing subassembly, and a supporting equipment subassembly. The processing subassembly comprises servers, field-programmable gate arrays (FPGA) hardware, FPGA signal processing code, and monitor and control code. The data routing subassembly comprises Programming Protocol-independent Packet Processors (P4) switch hardware, monitor and control code, and optical/copper communication. The supporting equipment subassembly consists of racks, power distribution, cooling, monitor and control networking, servers, and software. This design of the Correlator and (Array) Beamformer is called atomic commercial-off-the-shelf (COTS) design²⁹ (see Fig. 15).

The FPGA-based hardware has evolved over generations, starting with bespoke hardware, and now using COTS hardware, which significantly reduces the design and procurement time. The signal processing system design uses FPGA boards to provide the compute resources at low power and cost. The Xilinx Alveo cards were selected as a suitable choice for FPGA boards. They have integrated high bandwidth (HBM) memories capable of around 3.6 Tbps transfer rates. Communication is provided by a single 100 GbE quad small form factor pluggable (QSFP) per card.




Fig. 15 The subassemblies of the correlator and (array) beamformer.

Each FPGA card is air cooled and 20 of them are supported in a single dual-processor server chassis. External 100-GbE signals, both input and output, connect to the QSFP on the board front panel. Monitoring and control take place via a PCIe bus.

The Atomic COTS design philosophy is characterized by the following features:

1. All signal processing for all stations and a subset of the bandwidth is implemented in a single FPGA card. This is the “Atomic” processing.
2. All data routing is implemented externally to the FPGA card. This uses network switches programmed using the P4 language, which allows data to be routed on metadata within the packet, such as the sky frequency of the data in the packet.
3. Each FPGA card implements a single function such as correlation, Pulsar Search beamforming, or Pulsar Timing beamforming. This feature, in addition to the use of P4 switch, gives the system big flexibility and modularity. It greatly simplifies the firmware for each FPGA and decouples development on each function.

The FPGA cards are included in 21 servers, configured as 5 racks of 4 or 5 servers per rack, with 20 FPGA cards per server. All servers are identical and interchangeable. One of the servers and its FPGA cards act as an “engineering subarray,” used for diagnostic and maintenance, independent of the observing array. This extra server also allows any single server to be swapped out without affecting ongoing observations. The signal transport is a two-layer network with 9 I/O P4 switches connecting to external elements and 11 intermediate P4 switches with connections to the Alveo boards. The estimated power dissipation for the system is about 60 kW.

The Correlator and (Array) Beamformer is a highly configurable system supporting imaging, and non-imaging (Pulsar Search/Pulsar Timing) using different FPGA code images. In imaging mode, the Correlator and (Array) Beamformer supports multiple subarrays and station beams sharing the available 300-MHz bandwidth. Imaging parameters including delay polynomials, polarization corrections and beamformer weights must be updated regularly during an observation to maintain interferometric coherence.

The Signal Processing Subsystem delivers data from 384 coarse frequency channels per station and each 100 GbE link from Signal Processing Subsystem carries data for six stations. After the P4 data routing subassembly each FPGA card receives data from all stations for 2, 4, or 6 coarse channels depending on the FPGA card functionality. This is immediately buffered to HBM memory which greatly simplifies timing constraints on the Signal Processing Subsystem data flow. After buffering about 0.1 s of data, the data are delayed, corrected, and processed by a filter-bank. RFI detection then occurs on the fine channel data from the filter-banks and all downstream processing ignores RFI flagged data. This data is then buffered again to HBM and, when sufficient data is stored, processing of this data in the beamformer or correlator can take place.

For Pulsar Search, a station-based and output beam polarization Jones correction is applied. For Pulsar Timing, every input to each beam has a polarization correction. At the output both Pulsar Timing and Pulsar Search beams are aggregated across all P4 switch ports resulting in a single P4 output having full bandwidth data for three Pulsar Search beams or one Pulsar Timing beam. Beam data are routed through the P4 switch based on the beam number contained within the packet header.

For correlation, multiple fine channels are processed simultaneously to give the required frequency resolution for a subset of the visibilities. The variation in the number of fine channels allows standard 5.4-kHz resolution observing and zoom resolutions. The visibilities are buffered to HBM, and the next subset of visibilities calculated. When all data for a single output frequency visibility data are stored, the data set can be read out. As with the beamformer, packetization and readout to the data routing subassembly may be delayed to smooth the output data flow. Precision time protocol (PTP) system synchronization across all FPGA cards is used to control the delays and processing order. PTP timing permits to synchronize the processing of the received data from the stations so that no huge buffer of data is required and to avoid congestion on output links.

All data products, beams and visibilities, are output on the same 100 GbE link that provides station beam data as input to the FPGA cards. The two layers of the P4 switch network read the metadata in the packets and use this to route the data to the appropriate output.

The monitor and control software provides the TANGO interface to the LMC subsystem. A monitor and control server acts as a gateway for monitor and control communication over an internal 10 GbE network connecting with power, cooling, P4 switches, FPGA card servers, and a PTP switch.

6.4 Time-Domain Processing (Pulsar Search and Pulsar Timing Subsystems)

SKA1-LOW time-domain processing consists of two DSP subsystems. The PSS³⁰ is designed to search the entire sky uniformly for undiscovered pulsars as well as to discover rapid, single pulse-like signals known as Fast Transients (e.g., the recently discovered fast radio bursts). The PST subsystem³¹ is designed to precisely measure the TOA of known pulsars. Because of their distant origin, time-domain signals are characterized by dispersion, which causes the high-frequency components of the signals to arrive before the low-frequency components. This is caused by the pervasive existence of plasma along the path from the source to Earth and is an important means of distinguishing these signals from man-made ones.

6.4.1 Pulsar search subsystem

The PSS^{1,30} accepts up to 500 array beams from the Correlator and (Array) Beamformer and searches each in parallel for pulsars and transient sources. In the case of pulsars, the search space is potentially three-dimensional (3D): pulse period, Dispersion Measure, and if a pulsar is contained in an orbit around another mass, the rate of change of pulse period (acceleration). Each Pulsar Search processing node operates mostly independently on groups of three of the 500 array beams. The bandwidth of each beam is 118 MHz, channelized into 8192 channels.

The Pulsar Search Data Receptor receives the tied-array beamformed data, and the search is carried out by a process called de-dispersion, where the wide bands are divided into narrow frequency channels. In a search, the dispersion, parameterized by Dispersion Measure, is unknown and the search is made over a range of possible Dispersion Measure values. The signal in each channel is corrected for the delay caused by the candidate dispersion, and the channels can then be summed to reveal the pulse of the source. The number of dispersion measures to be processed is defined by the observing frequency and science requirements. The output of this process is a single de-dispersed time-series, for each of the trial dispersion measures.

The number of Dispersion Measure trials is selected such that de-dispersion can be done in real time while completing enough trials that the coarseness of sampling in Dispersion Measure phase-space, in the worst case, results in recovered signal-to-noise ratio (SNR) of <15% from the ideal SNR that would be measured at the exact Dispersion Measure of the signal. Source candidates are then sent to the SDP as a data product (Fig. 9).

The sections of the de-dispersed time-series are analyzed for the presence of transients, where bright single pulses in the time domain are detected. Pulses of different widths are detected by down-sampling in time, followed by applying a set of matched filters and then thresholding at a pre-set level above the noise. Single pulses are then sifted, where pulses of different widths and different Dispersion Measures belonging to the same astrophysical signal are grouped together as a single candidate. A fine search in dispersion measure and pulse width around the best Dispersion Measure and width is carried out for each candidate. The optimized candidates are also sent to the SDP as a data product (Fig. 9).

The PSS comprises high performance COTS computational hardware hosting bespoke software. Each Pulsar Search node consists of a CPU and accelerator hardware [FPGAs and graphic processing units (GPUs)] and can process several beams. All these components are controlled by the LMC subsystem.

6.4.2 Pulsar timing subsystem

The goal of Pulsar Timing is to measure the TOA (i.e., absolute time instant when a radio signal reaches a remote receiver) of the known pulses against UTC time, which is kept by an assembly of the world's atomic clocks, and to track the TOA of each pulsar for durations up to a decade.

As for transient detection, the procedure is also one of template-matching. But the process is significantly different from searching for pulsars because there is much more a priori information. This information is stored in a database that includes the properties of the template that are accumulated from previous measurements. In some cases, such as newly discovered pulsars, the information may not be very accurate; in other cases, it may represent years of observations.

The PST subsystem^{1,31} receives up to 16 dual-polarization array beams from Correlator and (Array) Beamformer, which are processed in parallel. A PST tied-array beam consists of channelized “voltage” (i.e., containing amplitude and phase) time series that are quantized at 16 bits per sample. It performs TOA processing on one pulsar per beam. The de-dispersion technique used for PST is more accurate but more compute-intensive than that used for Pulsar Search: phase-coherent de-dispersion that requires performing many large Fast Fourier Transform operations in real time.

PST implements DSP functions on enterprise-class servers equipped with high-speed networking and GPUs. The DSP architecture is constructed from modular functional pipelines, which can be selected, depending on the configured processing mode. In detail:

1. PST: coherent de-dispersion, RFI mitigation, Stokes detection, and phase-folding to generate Integrated Pulse Profiles. These are the frequency and pulsar phase resolved averages of polarized flux for a given pulsar.
2. Dynamic Spectrum: coherent de-dispersion, RFI mitigation, Stokes detection, and phase folding to generate high time, frequency, and polarization-resolved spectra re-quantized to efficient bit widths.
3. Flow Through: selection, re-quantization, and re-packing frequency channels to form an output voltage time series.

6.5 Science Data Processor

The SDP is responsible for the processing of observed data into science data products, the long-term preservation of these data products, and the delivery of these products to the SKA Observatory and the SKA Regional Centers. It is also responsible for contributing to the array-level calibration of the telescope and for providing quasi real-time quality metrics for ongoing observations.

SDP is a complex subsystem that provides various pipe-line processes that can be executed to calibrate the telescope and analyze the data. Figure 16 shows the architecture, the flow of data and the functions performed by the SDPs, a HPC. Figure 17 shows the many interfaces that SDP has with other subsystems. Under the control of the SKA1-LOW scheduling, monitor and control subsystems, SDP³²:

1. Ingests and buffers two types of observational data at a high rate (Fig. 16). These data are visibility data, received as a continuous flow to be processed and imaged, and time-domain data received as discrete chunks.
2. The visibility data are transformed into 3D images (cubes) containing two spatial and one frequency dimension. This processing is carried out in a highly parallelized manner as depicted in Fig. 16, both at the coarse scale (e.g., data for an entire frequency channel) and a fine scale (e.g., time and baseline division).
3. Time domain data are similarly processed in parallel. SDP receives pulsar candidates for further analysis, and these are persisted and processed in batch mode subsequently.
4. Output data products are stored in a science archive and also delivered to the SKA Regional Centers over intercontinental fiber networks.
5. Processing is executed in such a way as to maintain ingest availability. This requires considerable advance observation planning.

The main difference between the SKA1-LOW SDP and other generic HPC systems is that the SDP execution is scheduled as part of the SKA telescope observations so that the overall throughput of the processing does not result in observation blocking. The availability of data processing resources thus becomes a constraint in scheduling telescope observations, and the timely execution of data reduction functions is essential in maintaining the telescope operational.




Fig. 16 SDP dataflow, architecture and functions.




Fig. 17 SDP data exchanged by SDP with other telescope subsystems.

Recognizing that HPC hardware capabilities are rapidly evolving and the underlying software is very long-lived, the software architecture is designed with the following in mind:

1. Scalability and performance: ensure that the SDP system keeps up with the ongoing observations and data ingest while enabling the continuous processing of data in batch mode.

This allows it to address the wide range of scientific use cases enabled by SKA1-LOW. The raw required peak performance for SDP is of order 125 PFlop. The ingest rate of raw data into each system is of order 1 TB/s and the sustained read rate from persistent storage is of order 10 times this value. Scalability of the software architecture is essential so that an appropriately sized hardware system can be utilized. Different workflows will adopt different computing paradigms, ranging from embarrassing parallel processing to computationally expensive workflows.

2. **Modifiability:** SDP will support the development and execution of emerging workflows to meet new scientific challenges and use cases in the years to come. Due to the long lifetime of the SKA project, the SDP software will also have to be robust to hardware changes and evolution.
3. **Availability:** the SDP system is critical to the operations of the telescope and as such it needs to be highly robust and available. The large amount of data makes it prohibitively costly to store raw data as ingested by the SDP, thus imposing strict timing requirements on the processing system. This makes sure that computing resources are available when scheduled as part of an observation.

There are existing packages that individually meet one or a few of the above requirements but none meet all. The SDP architecture builds on the lessons learned in Common Astronomy Software Applications,^{32,33} Australian Square Kilometer Array Pathfinder (ASKAPSoft),³⁴ LOFAR,⁴ and MeerKAT³⁵ projects extensively by adopting an architecture explicitly designed to exploit the intrinsic data parallelism in the radio astronomy data and using where appropriate a data-flow approach to obtain scalability while supporting other large-scale programming models where these are better suited to the problem.

To achieve these goals the SDP system is organized into different loosely coupled software components, as described in the components and connectors view in Fig. 18.

The SDP ingest section receives data streams from the Correlator & (Array) Beamformer, Pulsar Search, Pulsar Timing and Monitor Control and Calibration subsystems. It performs immediate real-time processing and it stages the data into a large buffer. Offline processing




Fig. 18 SDP components and connectors view. LFAA indicates the monitor control and calibration subsystem; CSP the array correlator and beamformer, pulsar search and timing subsystems, SRC the SKA regional centers.

stages perform basic calibration, data reduction, and quality assessment. Multiple real-time and batch processing can execute simultaneously within the SDP and usually there will be at least one real-time and one batch instance always running.

The science data products obtained are stored in the “buffer” and delivered to the SKA Regional Centers via the “delivery” component. Data queues coordinate the internal real-time data exchange, which is also needed by quality assessment to feedback calibration information to the system. The execution control component is responsible for interfacing with Telescope Monitor and Control and to orchestrate processing as scheduled.

Execution frameworks and their instantiations as execution engines are key elements of real-time and batch processing. The implementation of execution frameworks within the SDP is therefore as follows (see Fig. 19):

1. A scheduled processing instance is described by a Processing Block (PB). The PB contains parameters needed for the processing to proceed and a reference to a workflow to execute.
2. A workflow is an algorithm to process data within the SKA. It is realized by a script/program, which: requests allocation of resources, instantiates multiple instances of execution frameworks (termed execution engines), manages the coarse distribution of data between these execution engines.
3. The architecture supports multiple execution frameworks: this is to account for different requirements of different processing stages.

This architectural model allows the SDP to implement a variety of different workflows,³² often described as science pipelines, in a series of processing steps, or tasks, with data flow between them. For those parts of the workflow for which a data driven approach is not the best solution, the SDP architecture permits the use of message-passing processing steps such as those realized by a message passing interface program. Workflows include, among others, Receive and Preprocessing, Calibration, Pulsar Search, and Pulsar Timing, Deconvolution, and Imaging.




Fig. 19 Batch processing primary representation.

The architecture allows new algorithms to be added simply by creating new scripts to realize them. Typically, a combination of various workflows is required to support an observing mode.

6.6 Observation Monitoring and Control System

The OMC system is the suite of software subsystems responsible for managing the planning, scheduling and execution of observing. This system is central to the overall SKA architecture, as it allows to correctly operate and monitor the observatory and the telescopes, and it constitutes the key interface between the science community users and the observatory. It also ensures efficient, complete, and reliable data and meta-data collection. The OMC is mainly constituted by two major subsystems: the OSO and the TMC.

The OSO is the suite of software tools that support the entire science process from proposal submission to data-product delivery. The OSO tools form a system of systems framed around a shared data model and use a shared data architecture to provide an easy-to-use human interface that hides the complexity of a complicated telescope as the SKA1-LOW from the users. A similar pattern is also in use by many other major world-class astronomical observatories, such as the Atacama Large Millimeter/submillimeter Array (ALMA),³⁹ Gemini,⁴⁰ and UKIRT.⁴¹

The interaction between the different OSO tools (see Fig. 20) allows the community investigators to submit observing proposals and have them turned into projects upon approval. Projects are then described as a set of atomic observing units called “Scheduling Blocks” (SB), encapsulating the technical details about the observation, and PBs encapsulating its processing requirements. Each SB will execute a number of scans, i.e., a period of observing where the telescope configuration stays the same while data are being recorded. Multiple integrations/scans fall within a SB. The observing activity of the SKA1-LOW Telescope is then organized by scheduling different SBs for execution at different times, along with the respective PBs. According to the schedule, SBs are executed by the Observation Execution Tool and the progress on each Project is monitored and reported to the users by the Project Tracking Tool (Fig. 20).

The TMC is the software system responsible for the lower-level instrumental monitoring and control of the telescope, and it performs the configuration, execution, and monitoring of observations. Its main responsibilities are identified as follows:

1. Support execution of astronomical observations.
2. Manage telescope hardware and software subsystems to perform astronomical observations.
3. Manage the data to support operators, maintainers, engineers, and science users to achieve their goals.
4. Determine telescope state.

The TMC is also responsible for keeping historical records of the measurements and attributes (monitor points), alarms, logs, hw/fw/sw versions, serial numbers, configurations, and more. These are stored in the Engineering Data Archive.

To maximize use of the SKA telescopes, telescope operators can subdivide the collecting area and processing resources into up to 16 subarrays, and operate each subarray as an independent telescope when it comes to: observing mode, observing targets, and start/stop time of the




Fig. 20 The different OSO tools implementing the proposal and execution flow.

observations. A subarray is a subdivision of the SKA1-LOW telescope that can be scheduled and operated independently of other subarrays. A subarray constitutes a set of resources (i.e., receptors, correlator slices, etc.) and can be as large as the whole telescope array, or a single constituent item. A subarray is only prevented from being created by resource constraints. Resources in this context include, but are not limited to, number of available stations, power availability, data transport bandwidth, signal processing, and data processing capacity.

An engineering subarray may not need an entire system slice, but this is not precluded. For instance, a set of stations without the need of signal or data processing may constitute an engineering subarray. Conversely, an astronomy subarray requires end-to-end capability, i.e., a full system slice.

A subarray is also the concept upon which all observing using SBs rests—a SB instance (SBI) requires a subarray for its execution. Each SB will include as meta-data a description of the resources it requires for execution, and these either directly or indirectly define the make-up of the subarray required.

The telescope control system plays a key role by providing the interface that allows operations to create and independently operate the subarrays.

To meet this complex set of requirements, the telescope control system is organized as a hierarchical set of control nodes: the telescope is subdivided in several large subsystems, according to the function and the technology used, and each subsystem is responsible for its own LMC, exposing a high-level Application Programming Interface (API). This enables separation of concerns: each subsystem provides a set of commands and exposes a set of resources and attributes that can be used by the entities higher in the telescope Monitor and Control hierarchy to monitor status, request state and mode changes, allocate resources, set observing modes, etc. The same approach is leveraged in monitoring and alarm reporting so that entities higher in the Monitor and Control hierarchy derive and report the overall status of a subsystem.³⁹

The TMC implementation is based on the TANGO⁴⁰ framework: an open-source device-oriented controls toolkit for building Supervisory Control and Data Acquisition systems in use in other large scientific facilities.⁴¹ The key concepts and features of the TANGO framework are:

- a. TANGO Device—a software model of a device (a piece of equipment or a subsystem). Each TANGO Device implements a standard API (commands, attributes, and properties).
- b. TANGO Device Server—a process that provides an execution environment for TANGO Devices.
- c. TANGO DB—a database where the TANGO Devices and Device Servers Register.
- d. TANGO bus—a communication bus which allows any TANGO Device to communicate with any other TANGO Device. It also allows TANGO clients to access the devices.
- e. TANGO Client—a software entity which uses the well-known TANGO Device name (fully qualified device name—FQDN) to access TANGO Devices, obtain read/write access to attributes, and issue commands. The FQDN can be obtained from the TANGO DB.

As part of the design process of the TMC, a number of prototyping activities have been carried on in order to evaluate the feasibility of the architecture in conjunction with the TANGO framework,³⁹ providing excellent results but also highlighting the necessity of developing standardized guidelines to harmonize and drive a standard approach to the implementation of the control and monitoring software, across the TMC and the various LMCs. As a result, a set of Control System Guidelines has been developed. These detail how each subsystem should layout a set of different TANGO devices to integrate within the TMC, and how different properties of each subsystem shall be exposed in the TANGO implementation. This results in an established set of nodes and patterns to be used when implementing the SKA1-LOW control system (see Fig. 21):

1. The Central Node is the top-level TANGO device in Monitor and Control hierarchy. It provides an aggregated view of the Telescope in terms of Telescope availability, utilization factor, and Health State. It derives telescope Health state from the health state of all the subsystems as reported by the respective Leaf Nodes. It is also responsible for subarray management, allocation/deallocation of resources; coordinating the control within the telescope for operations such as startup, shut down; and implementing operator instructions for critical and non-critical interventions.




Fig. 21 TMC LOW components and connector diagram describing the interaction between the main TANGO devices constituting the LOW control system. CSP includes the correlator and (array) beamformer, and the PST and PSSs (Fig. 9).

2. The Subarray Node groups all the resources that are participating in an observation into one logical entity. This allows uniform command and control for carrying out the observation. This ensures that observations are carried out independent to each other. It coordinates the configuration and execution of a scan as instructed directly from the Observation Execution Tool. It is also responsible for deriving and aggregating subarray level Health State. Central node aggregates physical health state from leaf nodes and observation health state from all the Subarray Nodes to derive the Telescope health state.
3. Leaf Nodes provide a very consistent way to communicate with all telescope subsystems and provide a layer managed by Telescope Monitor and Control, which can implement consistent rolled-up reporting and common behavior control. Leaf Nodes are responsible for communicating with the Element master (receiving monitoring data, Element status, and relaying commands). Execution of subsystem specific Algorithms or Engineering procedure, e.g., Pointing Model, Delay and phase models for beam forming, etc., is done at Leaf Node level. The Leaf Nodes translate the control command (commands received from Subarray Node) to corresponding subsystem specific commands. In case of a Leaf Node failure, it becomes easy to diagnose the failure and restart the respective Leaf Node. This is done in isolation without affecting the other leaf nodes. Elements such as the Central Signal Processor [i.e., Correlator and (Array) Beamformer, PST, and PSS as in Fig. 9] and SDP are shared across all the 16 subarrays, as such these Elements expose subarray specific controllers along with the Element Master device.

6.7 Synchronization and Timing System

To achieve the intended science goals of the SKA1-LOW telescope it is necessary to sample the astronomical data coherently captured by each receptor and tie it to the time as accurately as possible. Installing a high-quality reference frequency and time reference source at each remote location is not a practical solution. Instead of that, the Frequency and Time standards are maintained centrally (in the CPF) and are distributed to each RPF over a dedicated distribution system. The channel for distribution is fiber cable is exposed to environmental conditions such as diurnal temperature changes, which can introduce phase drifts in the signals, adversely impacting the time accuracy and the coherence of the reference frequency delivered. The SAT system maintains the time and frequency standards at the CPF, and distributes these standards to each RPF with automatic adjustments to maintain signal quality. Figure 22 shows a simplified block diagram of the SAT system, which consists of an observatory clock subsystem




Fig. 22 Simple block diagram of the SAT system used in the SKA1-LOW telescope.

and two separate distribution systems for time and frequency. For distributing time and frequency standards to the RPF, the SAT system uses dedicated distribution subsystems, which actively compensates for the delay variation due to environmental perturbations to maintain the desired quality of the references delivered. For those stations whose digitization happen in CPF, the SAT system distributes these references within CPF using distribution amplifiers.

The Observatory Clock supplies local time standards with an uncertainty better than 9 ns (1-sigma) reference to the UTC and a stable reference frequency to the distribution systems. The Observatory Clock subsystem uses three hydrogen masers connected in a “three-cornered hat” configuration. The masers in this configuration can be compared in pairs so that the incipient failure of an individual maser can be detected. This configuration guarantees redundancy by providing a master and backup realization of the clock. It is also a well-established technique⁴² for estimating the stability of the individual clocks. After comparing the performance of the clocks, a master clock is selected to supply the references to the rest of the distribution chain. The reference time signal at the output of the master clock is needed to be traceable to the UTC with the least possible uncertainty. Traceability to the UTC is typically achieved by Two-Way Satellite Time and Frequency Transfer (TWSTFT) or common view GNSS TFT. The TWSTFT method involves transmitting microwave Ku Band signals into space and is incompatible with use at the site of a radio telescope. Therefore, the SKA1-LOW will utilize common view GNSS TFT at the telescope site. GNSS “Common View” TFT make use of the received Clock Data from orbiting GNSS satellites and clock offset data is sent to another lab (BIPM in this case) for comparison. This technique allows direct comparison of the clocks located at remote locations, BIPM and SKAO in this case. BIPM will produce corrections after comparison, which then can be used to steer the clock to maintain the least uncertainty in the traceability to the UTC. This is possible by implementing SKAO’s own UTC(k) instance. In the traceability uncertainty, internal delays of the GNSS receives plays an important role. Hence the SAT system also includes the provision of an additional visiting/calibrating GNSS receiver to calibrate internal delays of the permanent GNSS receiver. The visiting GNSS receiver acquires the calibration from the BIPM (or from the BIPM approved lab) and transfer the calibration to the permanent GNSS receiver. This configuration works as local time and frequency standards that are maintained centrally. Apart from this, there are two different systems, one distributes time standards and the other distributes frequency standards. The distribution system for time standards utilizes the White Rabbit (WR) technology, which is an Ethernet-based protocol. The WR protocol is an extension of the PTP (IEEE1588), which offers better time accuracy over the network. The original WR system performance has been improved by incorporating specially designed Dense Wavelength Division Multiplexing filters and long-reach SFPs capable of running up to 80 km without the need for repeaters. It also tracks and keeps correcting the delay variations to deliver a time signal with an accuracy better than 2 ns compared to local time standards maintained at the CPF. The reference frequency distribution system measures and keeps track of the phase variations through the fiber and counteract to stabilize the phase fluctuations due to the fiber cable. It delivers the reference frequency to each RPF which is coherent across all the processing facilities. It delivers the reference frequency with coherence better than 98%. At the same time, it also cleans the short-term jitter at the endpoint using a jitter cleaning oscillator.




Fig. 24 NSDN topology. TM in this figure indicates the TMC.

and forwarding, and the flexibility to integrate security appliances such as firewalls at core locations when required.

The NSDN will use its own physical network to transport its own management signals in-band, and it will be managed by the network manager (NMGR) with a northbound interface to the TMC.

6.8.1 Logical and physical overview

Using network equipment and optical fiber cables, the networks provide data connectivity between many SKA1-LOW systems and subsystems in the three main locations in Australia: Boolardy, Geraldton, and Perth (Fig. 3).

Figure 25 shows the networks' high-level architecture between key telescope systems and the approximate fiber cable route distances between them. There will be optical fiber cables from the 12 RPFs along each of the three spiral arm routes back to the CPF. The cable distance from the furthest RPF to the CPF will be design limited to <80 km.

An existing fiber cable system with integrated optical amplifiers is already operational between the ASKAP control building, located in the Boolardy array region, and Perth via Geraldton. There are several repeater shelter huts on this route every ~100 km, that house the optical amplifier equipment. This existing fiber cable system is operated by the Australia's Academic and Research Network (AARNet) as part of the National Research and Education Network (NREN), and it will be extended from the ASKAP control building to the CPF. There will be an allocated number of fiber cores in the existing long-haul cable used to transport science and non-science data between the CPF and the Engineering Operation Center (EOC), and between the EOC and the SPC and Science Operation Center (SOC).




Fig. 25 Logical connectivity of the SKA1-LOW telescope.

The SKA1-LOW traffic will make use of most of this existing infrastructure and will be upgraded where necessary.

6.8.2 Network management

Management of the network elements of the telescope are primarily performed by the NMGR, which performs fault and performance management for the network elements and acts as the LMC for reporting status and performance information to TMC. It should be noted that other management systems exist to monitor performance of other systems within the telescope. The CPF-SPC link will be funded by the SKAO, but the Australian National Research Network AARNet will provide overall end-to-end operation, monitoring, and maintenance of the link by exchanging information with the NMGR servers. The NMGR servers report the network status, topology and performance of the CPF-SPC link to the TMC, through the NSDN.

7 Conclusion

The system CDR for SKA1 was completed in 2020, endorsing the design, project plans, and budgets to start the construction, which has now begun. The SKA Observatory Council has endorsed the SKA Construction Plan,¹⁹ which incorporates the detailed project plans and budgets to construct SKA1 according to the approved design. This paper has shared highlights of the SKA1-LOW architecture and subsystem design, showing how the telescope addresses some of the technical challenges faced by an ambitious facility such as the SKA. The resulting design implementations have been described, including the latest design improvements,^{1,29} on which further detail can be found in the references provided.^{16,21,26,27,30–32,42–46}

Over and above the system and subsystem design work, there has been a significant amount of pathfinder and pre-cursor development,^{4,6,35} performance modeling,⁴⁷ and hardware and software prototyping and testing^{18,29,48–50} to confirm the feasibility and cost of the presented designs. The SKA1-LOW 256-antenna aperture array and station processing has undergone two prototyping iterations on site,^{49,51} and a separate prototype system integration facility (PSI) has been developed off site (CSIRO facilities, Marsfield, Australia) to house a partial implementation of

the SKA1-LOW signal chain and hosting the early deployment of the SKA software system and data analysis components. With the kickoff of construction activity, a renewed focus on final characterization of the evolved designs as production quality prototypes will see further development of these verification installations to demonstrate the completion of the SKA1-LOW design phase and usher in the initial six station Array Assembly planned for 2023.

SKA1-LOW has been through a number of independent reviews and has now reached the construction stage, inspiring confidence that the design will meet the performance, operational and cost challenges of the SKA community.

Acknowledgments

The authors wish to acknowledge the consortia/collaborating institutions teams and the SKAO engineering, software and science teams that have jointly contributed to the design and architecture development reported in this paper. The Project respectfully acknowledges the Wadjuri people of Western Australia as traditional owners and custodians of the land on which the SKA1-LOW telescope will be constructed. The authors have no conflicting interests related to this paper.

References

1. P. E. Dewdney et al., “SKA1 design baseline description,” SKA document SKA-TEL-SKO-0001075 (2019).
2. R. Braun et al., “Advancing astrophysics with the Square Kilometre Array,” in *Proc. AASKA14* (2014).
3. M. Caiazzo, “SKA phase 1 system requirements specification,” SKA document SKA-TEL-SKO-0000008 (2017).
4. M. P. van Haarlem et al., “LOFAR: the LOW-Frequency ARray,” *Astron. Astrophys.* **556**, A2 (2013).
5. G. Taylor, “First light for the first station of the long wavelength array,” *J. Astron. Instrum.* (2012).
6. S. W. Ellingson et al., “The long wavelength array,” *Proc. IEEE* **97**(8), 1421–1430 (2009).
7. S. J. Tingay et al., “The Murchison widefield array: the Square Kilometre Array precursor at low radio frequencies,” *Publ. Astron. Soc. Aust.* **30**, E007 (2013).
8. W. van Cappellen et al., “LOFAR2.0,” (2021).
9. W. van Cappellen et al., “LOFAR2.0,” URSI GASS (2021).
10. P. Dewdney, “SKA1 system baseline design v2,” SKA document SKA-TEL-SKO-0000002 (2016).
11. R. Braun et al., “Anticipated SKA1 science performance,” SKA document number SKA-TEL-SKO-0000818 (2017).
12. R. A. Laing, “Calibration plan,” SKA document SKA-TEL-SKO-0001640 (2019).
13. B. Wang et al., “Square Kilometer Array telescope —precision reference frequency synchronisation via 1f-2f dissemination,” *Sci. Rep.* **5**, 13851 (2015).
14. P. Boven, “DWDM stabilized optics for white rabbit,” in *Eur. Freq. and Time Forum*, pp. 213–216 (2018).
15. A. J. Otto, F. Di Vruno, and T. Nkawu, “SKA EMI/EMC standards, related procedures and guidelines,” SKA document SKA-TEL-SKO-0000202 (2021).
16. SKAO Science Team, “SKA1 level 0 science requirements,” SKA document SKA-TEL-SKO-0000007 (2015).
17. P. Bolli et al., “Test-driven design of an active dual-polarized log-periodic antenna for the Square Kilometre Array,” *IEEE Open J. Antennas Propag.* **1**, 1253–1263 (2020).
18. G. Macario et al., “Characterization of the SKA-low prototype station aperture array verification system 2,” *J. Astron. Telesc. Instrum. Syst.* **8**, 011014 (2022).
19. J. McMullen et al., “SKA phase 1 construction proposal,” SKA document SKA-TEL-SKO-1105 (2021).
20. P. Dewdney et al., “SKA1 low configuration coordinates—complete set,” SKA document number SKA-TEL-SKO-0000422 (2021).

21. F. Perini et al., “RF over fibre technology for SKA low receiver,” *J. Astron. Telesc. Instrum. Syst.* **8**, 011016 (2022).
22. P. Bolli et al., “Computational electromagnetics for the SKA-low prototype station AAVS2,” *J. Astron. Telesc. Instrum. Syst.* **8**, 011017 (2022).
23. E. de Lera Acedo et al., “SKALA, a log-periodic array antenna for the SKA-low instrument: design, simulations, tests and system considerations,” *Exp. Astron.* **39**, 567–594 (2015).
24. E. de Lera Acedo et al., “Evolution of SKALA (SKALA-2), the log periodic array antenna for the SKA-low instrument,” in *Int. Conf. Electromagn. in Adv. Appl.* (2015).
25. E. de Lera Acedo, H. Pienaar, and N. Fagnoni, “Antenna design for the SKA1-low and HERA super radio telescopes,” in *Int. Conf. Electromagn. Adv. Appl.* (2018).
26. F. Schillirò et al., “Design and prototyping of the Italian Tile Processing Module 1.6 (ITPM) for the low-frequency aperture array deployment,” *Proc. SPIE* **11445**, 1144574 (2020).
27. A. Magro et al., “A monitoring and control prototype for the SKA low frequency aperture array,” in *Int. Conf. Control, Autom. and Diagnosis* (2017).
28. A. Magro et al., “A software infrastructure for firmware-software interaction: the case of TPMs,” in *Int. Conf. Signals and Syst.* (2017).
29. G. Hampson et al., “Square kilometre array low atomic commercial off-the-shelf correlator and beamformer,” *J. Astron. Telesc. Instrum. Syst.* **8**(1), 011018 (2022).
30. L. Levin et al., “Pulsar searches with the SKA,” in *Proc. IAU Symp. 337: Pulsar Astrophys. - The Next 50 Years* (2017).
31. B. W. Stappers et al., “The prospects of pulsar timing with new-generation radio telescopes and the Square Kilometre Array,” *Philos. Trans. R. Soc. A* (2018).
32. P. Alexander et al., “SKA1 SDP high level overview,” SKA document SKA-TEL-SDP-0000180 (2019).
33. J. P. McMullin et al., “CASA Architecture and Applications,” *Astronomical Data Analysis Software and Systems XVI*, ASP Conference Series, R. A. Shaw, F. Hill, and D. J. Bell, Eds., Vol. 376, p. 127, San Francisco (2007).
34. T. Cornwell et al., “ASKAP science processing,” ASKAP-SW-0020, 2016, <https://www.atnf.csiro.au/projects/askap/ASKAP-SW-0020.pdf>.
35. SKA Observatory, “Precursors and pathfinders,” 2021, <https://www.skatelescope.org/precursors-pathfinders-design-studies/>.
36. J. Schwarz, A. Farris, and H. Sommer, “The ALMA software architecture,” *Proc. SPIE* **5496**, 190–204 (2004).
37. S. Wampler, “Software design of the Gemini 8-m Telescopes,” *Proc. SPIE* **2871** (1997).
38. A. Bridger et al., “ORAC: a modern observing system for UKIRT,” *Proc. SPIE* **4009**, 227–238 (2000).
39. V. Mohile, “SKA1 TM prototyping report,” SKA document SKA-TEL-SKO-0000148 (2018).
40. Tango Controls, “Tango controls software community,” <https://www.tango-controls.org> (2015).
41. A. Götz et al., “State of the Tango controls kernel development in 2019,” in *Proc. ICALEPCS’19*, New York, pp. 1234–1239 (2019).
42. A. S. Hendre et al., “Precise timescale, frequency and time-transfer technology for the Square Kilometer Array,” *J. Astron. Telesc. Instrum. Syst.* **8**, 011022 (2022).
43. S. Lloyd, “Low end to end architecture diagram,” SKA document number SKA-TEL-SKO-0001077 (2017).
44. S. Amy and J. Muller, “Non-science data network system design report,” SKA document number SKA-TEL-SaDT-0000301 (2016).
45. P. Alexander et al., “SDP science pipeline workflow view,” SKA document SKA-TEL-SDP-0000013 (2019).
46. S. R. Chaudhuri, “TMC software architecture document,” SKA document SKA-TEL-TM-0000242 (2018).
47. B. Alachkar et al., “Assessment of the performance impact of direction-independent effects in the Square Kilometre Array,” *J. Astron. Telesc. Instrum. Syst.* **8**, 011020 (2022).

48. R. Wayth et al., “The engineering development array 2: design, performance and lessons from an SKA-low prototype station,” *J. Astron. Telesc. Instrum. Syst.* **8**, 011010 (2022).
49. P. Benthem et al., “The aperture array verification system 1: system overview and early commissioning results,” *Astron. Astrophys.* **655**, A5 (2021).
50. R. Brederode, A. Pellegrini, and L. Stringhetti, “*SKA1 prototyping report*,” SKA document SKA-TEL-SKO-0001631 (2019).
51. A. J. J. van Es et al., “A prototype model for evaluating SKA-LOW station calibration,” *Proc. SPIE* **11445**, 1144589 (2020).

Maria G. Labate is the telescope engineer for the SKA1-LOW telescope at the SKA Observatory. She received her BS and MS degrees in telecommunication engineering from Mediterranea University of Reggio Calabria in 2005 and 2007, respectively, and her PhD in electronic engineering from Second University of Naples in 2008. She is the author of more than 40 journal and conference papers and co-inventor of three patents. Her current role is to lead the technical activities of the SKA Low-frequency telescope. She is a member of INCOSE.

Biographies of the other authors are not available.